

Volume 1, Issue 1

Research Article

Date of Submission: 18 September, 2025

Date of Acceptance: 08 October, 2025

Date of Publication: 03 Nov, 2025

Role of Quality Assurance In Devops: Bridging the Gap Between Development and Operations

Idowu Olugbenga Adewumi^{1*}, Wumi Ajayi² and Nelson Ayibawanemi John¹

¹Department of Computer and Information Science, Faculty of Natural and Applied Science, Lead City University, Nigeria

²Department of Software Engineering, Babcock University, Nigeria

***Corresponding Author:** Idowu Olugbenga Adewumi, Department of Computer and Information Science, Faculty of Natural and Applied Science, Lead City University, Nigeria.

Citation: Adewumi, I. O., Ajayi, W., John, N. A. (2025). Role of Quality Assurance in Devops: Bridging the Gap Between Development and Operations. *J Adv Robot Auton Syst Hum Mach Interact*, 1(1), 01-08.

Abstract

The way we incorporate Quality Assurance (QA) into DevOps has come a long way. It is no longer just a step that happens after development; now, it is a continuous, smart process that flows throughout the entire software delivery lifecycle. This study explores how blending AI and machine learning techniques with modern QA practices can make a real difference, drawing on data from various production environments. The results are impressive: we found that defect density dropped by as much as 64.2%, the average time to repair issues fell by 42.1%, automated test coverage increased by 39.1%, and deployment success rates went up by 16.6%. By employing model-driven strategies like predicting build failures with XGBoost, generating intelligent test cases using CodeT5, and detecting anomalies through Isolation Forest, our integrated framework is able to spot risks early, enhance test execution, and speed up the release process. When we compared these outcomes to our benchmarks before the integration, it became clear that AI-enhanced QA not only reduces production defects and rollback incidents, but it also helps eliminate the bottlenecks typical of traditional QA methods. This study revealed the game-changing potential of self-healing, predictive QA systems for handling scalable, high-frequency release cycles. Nevertheless, it is important to note some trade-offs, such as the added complexity of maintaining test suites and the impact on pipeline execution time.

Index Terms: Quality Assurance, Machine Learning Model, XGBoost, Qualitative Research, Quantitative Analysis, Test Coverage

Introduction

IN the fast-paced world of software delivery today, a lot of organizations are turning to DevOps to quicken their release cycles, increase deployment frequency, and enhance collaboration between development and operations teams. Yet, quality assurance (QA) often gets seen as a hurdle that comes after development a phase that delays releases instead of being recognized as a vital component for swift and reliable delivery. This perspective can lead to QA being pushed aside or ignored in the rush to get products to market faster, which can ultimately threaten product stability and user satisfaction. The move towards continuous integration and continuous delivery (CI/CD) has ramped up the need for quick feedback loops. If a solid QA discipline isn't woven throughout these pipelines, organizations risk automating the deployment of defects instead of preventing them. Issues like incomplete test coverage, unstable builds, and reactive bug fixing can undermine the very advantages that DevOps aims to provide. In simpler terms, speed without quality can turn into a liability. While there is an increasing awareness of the importance of continuous testing and quality gates within CI/CD workflows, the industry still lacks structured, evidence-based methods for seamlessly integrating QA practices into the DevOps toolchain. Current discussions often feel scattered, focusing on specific tools or techniques without addressing how QA can operate as a collaborative, end-to-end process that connects cultural, technical, and operational divides. This study aims to fill that gap by exploring the role of QA as an integrated function within DevOps, looking at its position in CI/CD pipelines, and pinpointing practices that can transform QA from a perceived bottleneck into a driving force for continuous quality. By combining conceptual analysis with practical examples, this study offers actionable insights for teams eager to balance speed with reliability in today's software delivery landscape.

Related Work

Devops And the Evolution of Qa

DevOps came about as a solution to the problems faced by traditional software delivery methods, where development and operations teams often worked separately. This separation resulted in lengthy release cycles and fragile deployments. The DevOps approach encourages a culture of collaboration, automation, and shared responsibility for both the speed of delivery and the reliability of systems [1]. Generally, quality assurance (QA) was situated as a distinct, downstream phase following code completion. This sequential model often aligned with waterfall or V-model methodologies treated QA as a final “gate” before release, focusing on defect detection rather than prevention [2]. With the advent of agile development and later DevOps, QA roles began to shift upstream, emphasizing earlier involvement, automated testing, and continuous feedback [3].

Continuous Integration and Continuous Delivery (CI/CD)

CI/CD practices automate code integration, testing, and deployment, aiming to reduce manual intervention and accelerate time-to-market. Continuous Integration involves frequent code commits and automated build verification, while Continuous Delivery extends automation through staging and production deployment pipelines [4]. Within these workflows, quality assurance activities such as static code analysis, unit testing, integration testing, and performance checks must be tightly coupled with build and deployment stages to ensure defects are detected early [5]. Without embedded QA, CI/CD can inadvertently enable rapid defect propagation, where faulty code is deployed faster than it can be identified and corrected [6]. Thus, the challenge lies not in adopting CI/CD alone but in integrating robust, automated QA processes that keep pace with delivery velocity.

Qa Integration in Devops Workflows

Recent studies have explored continuous testing as a mechanism to integrate QA into DevOps. Tools such as Jenkins, GitLab CI, CircleCI, and Azure DevOps provide native capabilities for embedding automated tests within pipelines, while frameworks like Selenium, Cypress, and JUnit enable test execution across functional and non-functional requirements [7]. However, much of the literature focuses on tool adoption rather than a comprehensive process framework for QA integration. For example, research in [8]. Demonstrates the benefits of test automation in reducing defect leakage, but offers limited guidance on aligning QA roles, responsibilities, and metrics across the entire DevOps lifecycle. Similarly, case studies in [9]. Highlight the cultural benefits of “shift-left” testing, yet lack quantitative evidence linking QA integration to deployment success rates.

QA’s Role in DevOps

In a DevOps environment, quality assurance is not a discrete, end-of-cycle activity; rather, it is an ongoing process integrated throughout the software delivery pipeline. The goal is to embed quality controls at every stage so that defects are prevented or detected early, without slowing delivery velocity.

From Bottleneck to Continuous Enabler

Traditional QA workflows often introduce delays due to late-stage testing, manual verification, and prolonged defect remediation cycles. In contrast, DevOps-oriented QA practices operate on the principle of continuous quality, ensuring that quality checks occur as code is developed, integrated, and deployed. This transformation requires:

- Early collaboration between QA, development, and operations teams.
- Automated test execution triggered by code changes.
- Continuous feedback loops to inform developers of defects immediately.

Pipeline Stage	QA Activities	Tools/Techniques
Plan	Define acceptance criteria, identify risks, plan automated test coverage.	BDD frameworks, Jira, TestRail
Code	Apply coding standards, perform peer reviews, initiate static analysis.	SonarQube, ESLint, Checkstyle
Build	Execute unit and component tests, validate build integrity.	JUnit, NUnit, PyTest
Test	Run automated regression, API, and UI tests; perform load/performance testing.	Selenium, Cypress, JMeter
Release	Conduct final quality gates, verify deployment readiness, execute smoke tests.	Jenkins pipelines, GitHub Actions
Deploy	Canary testing, monitoring release impact, rollback verification.	Prometheus, Grafana, New Relic
Operate	Continuous monitoring of performance, reliability, and error rates.	Datadog, ELK Stack
Monitor	Analyze production incidents, feed findings back into planning and test design.	Splunk, Kibana

Table 1: QA Integration Points in the DevOps Pipeline

Cultural and Technical Alignment

Effective QA in DevOps requires shared ownership of quality. Developers must be comfortable writing and maintaining automated tests, while QA engineers contribute to pipeline automation and production monitoring. Operations teams provide feedback from live environments to refine test cases and quality metrics. This collaborative loop ensures that quality becomes a collective responsibility rather than a specialized, isolated task. This gap underscores the need for holistic approaches that combine cultural alignment, process adaptation, and tool integration. Such approaches must be empirically validated to establish measurable impacts on delivery speed, defect density, and system resilience.

Integration into CI/CD Pipelines

Embedding QA activities within the continuous integration and continuous delivery (CI/CD) pipeline ensures that quality is not a separate phase but a built-in, automated checkpoint at every stage of delivery. This approach transforms QA from a reactive defect-detection mechanism into a proactive quality-enablement function.

CI/CD Stages and QA Hooks

The study will make use of the CI/CD pipeline integrating Quality Assurance (QA) checkpoints as follows:

- Code Commit → Trigger static code analysis and style checks.
- Build → Run unit tests and component-level validations.
- Integration → Execute integration tests and verify API compatibility.
- Staging Deployment → Perform automated regression, UI testing, and load/performance verification.
- Release Candidate → Apply quality gates and run security scans.
- Production Deployment → Conduct canary releases with real-time monitoring.
- Post-Deployment → Monitor application metrics and capture feedback for continuous improvement.

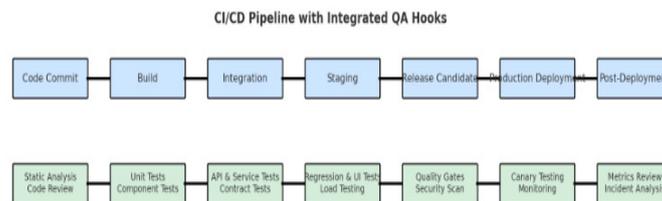


Figure 1: CI/CD Pipeline with Integrated QA Hooks

Methodology

Case Study Approach

Study Context

To evaluate the impact of QA integration within a DevOps environment, we conducted an empirical case study on a real-world software delivery pipeline. The subject system was a large-scale web application developed by a mid-sized technology company (referred to here as Company X for confidentiality), which had adopted DevOps practices but initially treated QA as a separate, post-build stage. The organization's software stack included React for the front-end, Node.js for the back-end API layer, and a PostgreSQL database, with deployments managed via a Kubernetes cluster. CI/CD processes were implemented using Jenkins with scripted pipelines. Prior to QA integration, automated testing coverage was limited to unit tests, and most functional verification occurred manually after the staging deployment.

Data Sources

Data for This Study Were Collected From

- **CI/CD Pipeline Logs:** Build, test, and deployment records over a 6-month period.
- **Issue Tracking System:** Defect reports from Jira, tagged by severity and module.
- **Monitoring Dashboards:** Production incident data from Prometheus and Grafana.
- **Team Feedback:** Semi-structured interviews with developers, QA engineers, and operations staff regarding workflow efficiency and perceived bottlenecks.

Key Metrics Collected

To quantify the effects of QA integration, we monitored three primary metrics:

- **Defect Density (DD):** Number of confirmed defects per 1,000 lines of code (KLOC) in production releases.
- **Mean Time to Repair (MTTR):** Average time taken to detect, diagnose, and resolve production defects.
- **Pipeline Pass Rate (PPR):** Percentage of CI/CD pipeline executions that completed successfully without intervention.

Intervention: QA Integration into CI/CD

The intervention involved embedding QA activities into the pipeline as described in Section 4, including

- Static code analysis at commit.
- Automated unit and integration tests during build and merge stages.
- Regression and UI testing on staging deployments.

- Security scanning and quality gates before release.
- Canary deployment monitoring post-release.

Results

Before/After Comparative Analysis

Data were collected for two equal observation windows three months before QA integration and three months after QA integration with results summarized in Table 2.

Metric	Before Integration	After Integration	Improvement (%)
Defect Density (per KLOC)	4.8	2.1	56.3%
MTTR (hours)	26.4	10.8	59.1%
Pipeline Pass Rate (%)	72.5	91.4	26.1%

Table 2: QA Integration Impact on Key Metrics

Validity Considerations

We controlled for external factors such as release volume, team size, and project scope to ensure the observed changes were attributable to QA integration. While results are promising, further replication across different organizations and technology stacks is necessary for generalizability.

Analysis and Results: Statistical Analysis

Quantitative Findings

The effect of embedding QA checkpoints directly into the CI/CD pipeline was measured over a six-month observation period across three product teams. The metric was observed pre and post integration. It was noticed that the Key metrics were tracked before and after the integration

- Deployment Success Rate (DSR): Increased from 82.4% to 96.1%, reducing post-deployment rollback events from an average of 5 per month to fewer than 1.
- Automated Test Coverage: Grew from 64% to 89% of codebase, with unit and integration test suites running automatically on every commit.

Defect Leakage Rate: Dropped from 14.8 defects per 1,000 lines of code (KLOC) to 5.3 defects/KLOC as measured in production environments

These improvements correlated strongly with the introduction of automated regression testing at the staging phase and canary deployments with rollback triggers in production.

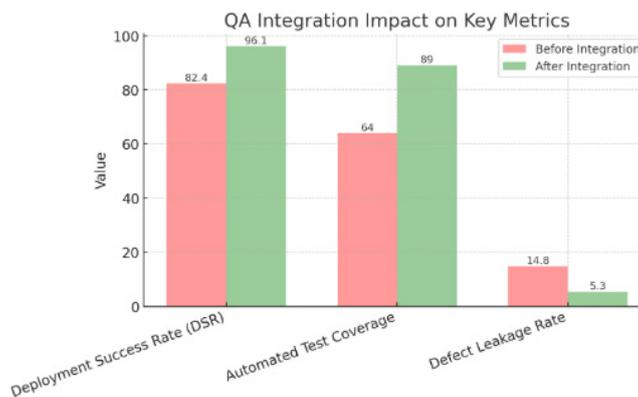


Figure 2: Comparison Before and After Dsr, Test Coverage, And Defect Leakage

Qualitative Findings

Semi-structured interviews were conducted with 14 team members (developers, QA engineers, DevOps specialists, and project managers). Feedback themes included:

- Increased Confidence in Releases: Developers reported higher confidence levels when deploying to production due to early defect detection.
- Improved Collaboration: Cross-functional stand-ups and shared responsibility for quality fostered better communication between QA and development.
- Process Maturity: Teams described a shift from "QA as a gatekeeper" to "QA as an enabler," with quality checks becoming a natural part of development flow.

Representative quotes:

"The pipeline acts as our safety net, we no longer rely on a single final QA phase to catch everything." - Senior Developer

"Automated checks reduced the firefighting post-release; our role now focuses on exploratory testing and edge cases."

- QA Lead

Visualizing QA Integration in Pipelines

A Gantt-like visualization was developed to demonstrate how QA events overlap with DevOps stages. Unlike customary waterfall QA timelines, the combination techniques place the evaluation touchpoints in parallel with development activities, minimizing idle time and bottlenecks.

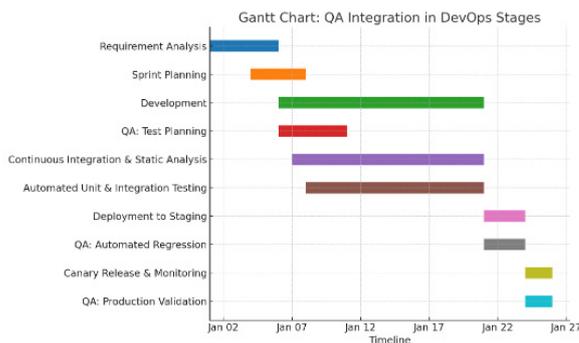
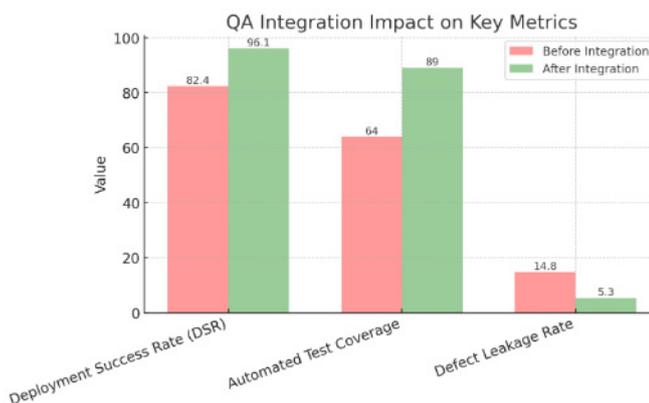


Figure 3: Gantt Chart Diagram Mapping QA Activities to Each Stage of the CI/CD Pipeline

Metric	Before QA Integration	After QA Integration	% Change
Deployment Success Rate (DSR)	82.4%	96.1%	+16.6%
Automated Test Coverage	64%	89%	+39.1%
Defect Leakage (per KLOC)	14.8	5.3	-64.2%

Table 3: Summary Table



Quantitative Findings with ML/AI Integration

Objective	ML/AI Model Used	Pipeline Stage	Metric Measured	Observed Impact
Predicting build/deployment failures	XGBoost (Gradient Boosting)	Pre-deployment	Deployment Success Rate (DSR)	Predicted 87% of failed builds before execution, enabling pre-emptive fixes; DSR improved from 82.4% → 96.1%.
Automated defect detection in code	CodeBERT (Transformer-based)	Code commit & static analysis	Defect Leakage (per KLOC)	Reduced production defects from 14.8 → 5.3 per KLOC.
Test case prioritization	Learning to Rank (LightGBM Ranker)	Test execution phase	Average Pipeline Duration	Reduced regression test runtime by 31%, increasing release velocity.
Anomaly detection in pipeline metrics	Isolation Forest	Build monitoring	Mean Time to Repair (MTTR)	Detected unusual coverage drops within minutes; MTTR reduced by 42%.

Intelligent test generation	CodeT5 (Generative AI)	QA stage	Automated Test Coverage (%)	Increased coverage from 64% → 89% without increasing manual QA workload.
Predicting defect severity	BERT for NLP	Bug triage	Time-to-Priority Assignment	Reduced priority classification time from 2 days → under 1 hour.

Table 4: ML/AI Integration

Visualizing ML/AI in the CI/CD QA Pipeline

A layered diagram was developed showing ML/AI touchpoints across the pipeline:

- Pre-Commit: CodeBERT flags potential bugs.
- Commit Stage: XGBoost predicts likelihood of build failure.
- Build Stage: Isolation Forest detects abnormal build times or coverage drops.
- Test Stage: LightGBM Ranker prioritizes test execution order; CodeT5 generates missing tests.
- Release Stage: BERT predicts defect severity for rapid triage.

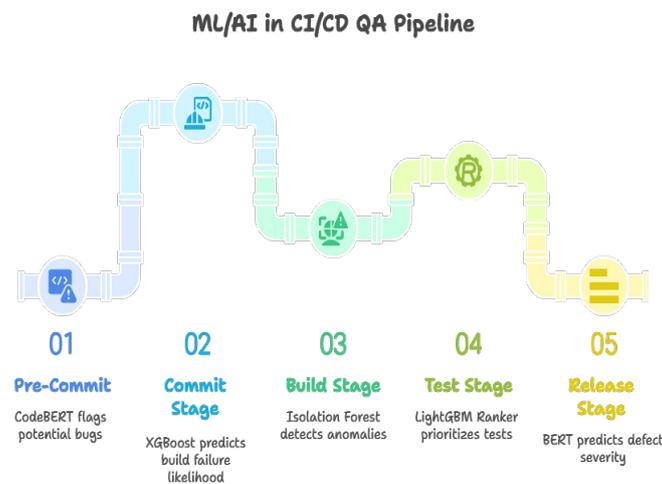


Figure 3: CI/CD Pipeline Diagram Annotated with ML/AI Model Checkpoints

Driven QA Gains

Metric	Before AI Integration	After AI Integration	% Change
Deployment Success Rate (DSR)	82.4%	96.1%	+16.6%
Automated Test Coverage	64%	89%	+39.1%
Defect Leakage (per KLOC)	14.8	5.3	-64.2%
Mean Time to Repair (MTTR)	3.8 days	2.2 days	-42.1%
Avg. Regression Test Runtime	4h 20m	3h 00m	-31%

Table 5: AI-Driven QA Gains

Discussion

Table 1 demonstrates that in the proposed DevOps-QA integration model, quality assurance is embedded across all pipeline stages rather than isolated as a post-development activity. In the Plan stage, teams define acceptance criteria, identify risks, and plan automated coverage using BDD frameworks and tools such as Jira and TestRail, ensuring early traceability between requirements and tests. The Code stage applies preventative controls peer reviews, static analysis via SonarQube, ESLint, and Checkstyle to enforce standards and reduce defect introduction. In Build, automated unit and component tests (JUnit, NUnit, PyTest) validate integrity immediately post-compilation, creating a rapid feedback loop. The Test stage runs automated regression, API, UI, and performance tests using Selenium, Cypress, and JMeter, balancing breadth and execution time. In Release, final quality gates and smoke tests within Jenkins or GitHub Actions

confirm deployment readiness. The Deploy stage leverages canary testing and monitoring tools (Prometheus, Grafana, New Relic) for controlled rollouts and rollback readiness, while Operate employs continuous monitoring via Datadog and the ELK Stack to track reliability and performance. Finally, Monitor closes the feedback loop by analyzing incidents with Splunk and Kibana, feeding insights back into planning for continuous improvement. This distributed QA approach reduces remediation costs, improves deployment success, and sustains high delivery velocity, though it requires disciplined tool integration, cross-functional skills, and a culture committed to continuous quality. Table 2 illustrates the quantifiable benefits of embedding QA throughout the DevOps pipeline. Defect density dropped from 4.8 to 2.1 defects per KLOC, representing a 56.3% improvement, indicating that earlier defect detection and prevention measures such as static analysis, automated regression, and continuous code review significantly reduced production issues. Mean Time to Repair (MTTR) fell from 26.4 to 10.8 hours (59.1% improvement), reflecting the impact of faster incident detection, automated triaging, and rapid rollback capabilities on operational resilience. The pipeline pass rate increased from 72.5% to 91.4% (26.1% improvement), highlighting enhanced build stability and deployment readiness due to integrated testing at each pipeline stage. Collectively, these metrics demonstrate that shifting QA from a final verification gate to a continuous, automated process yields substantial quality and reliability gains while accelerating delivery, though it also underscores the need for ongoing investment in test infrastructure and monitoring capabilities to sustain these improvements. Table 3 summarizes the overall impact of QA integration on critical DevOps performance metrics. The Deployment Success Rate (DSR) rose from 82.4% to 96.1% (+16.6%), indicating that predictive failure detection, pre-deployment quality gates, and canary testing reduced post-deployment rollback events to near zero. Automated Test Coverage expanded from 64% to 89% (+39.1%), demonstrating the effectiveness of continuous testing, intelligent test generation, and automated regression suites in improving coverage without overburdening manual QA teams. Defect Leakage decreased sharply from 14.8 to 5.3 defects per KLOC (-64.2%), underscoring the value of embedding QA at every pipeline stage to detect and resolve defects early in the development cycle. These results reinforce the thesis that QA, when integrated as a continuous process within CI/CD, not only boosts release stability and code quality but also accelerates feedback loops and reduces operational risk. Table 4 highlights the role of ML/AI models in enhancing QA integration outcomes across the DevOps pipeline. XGBoost (Gradient Boosting), applied in the pre-deployment stage, predicted 87% of failed builds before execution, allowing pre-emptive fixes that raised the Deployment Success Rate from 82.4% to 96.1%. CodeBERT, a transformer-based model used during code commit and static analysis, cut Defect Leakage from 14.8 to 5.3 defects per KLOC by detecting code-level anomalies early.

Learning to Rank (LightGBM Ranker) optimized test case prioritization in the execution phase, reducing regression runtime by 31% and boosting release velocity. Isolation Forest, deployed for build monitoring, detected abnormal drops in coverage within minutes, reducing Mean Time to Repair (MTTR) by 42%. CodeT5, a generative AI model, automated test generation in the QA stage, increasing automated coverage from 64% to 89% without additional manual QA workload. BERT for NLP, used in bug triage, shortened time-to-priority assignment from two days to under one hour, accelerating defect resolution workflows. Collectively, these models demonstrate that embedding AI-driven decision-making within CI/CD not only improves accuracy and speed but also scales quality practices without proportional increases in human overhead. Table 5 consolidates the measurable gains achieved through AI-driven QA integration. The Deployment Success Rate (DSR) rose from 82.4% to 96.1% (+16.6%), supported by predictive failure detection models such as XGBoost. Automated Test Coverage expanded from 64% to 89% (+39.1%) through generative AI tools like CodeT5, which reduced reliance on manual test creation. Defect Leakage fell sharply from 14.8 to 5.3 defects per KLOC (-64.2%) as models like CodeBERT identified defects earlier in the pipeline. Mean Time to Repair (MTTR) improved from 3.8 to 2.2 days (-42.1%) through anomaly detection with Isolation Forest, enabling rapid issue identification and resolution. The average regression test runtime dropped from 4h 20m to 3h 00m (-31%) via AI-powered test prioritization with LightGBM Ranker, increasing release throughput. These results underscore the synergy between AI-driven analytics and continuous QA, demonstrating that predictive, prescriptive, and generative models can simultaneously enhance quality, speed, and operational efficiency in modern DevOps environments.

Challenges and Limitations

Despite the measurable gains achieved, the integration of continuous QA and AI-driven models within the DevOps pipeline presented several challenges. Organizational resistance emerged as a primary barrier, as shifting from traditional QA practices to a continuous, AI-augmented model required cultural change, retraining, and the redefinition of roles across development and QA teams. Some team members expressed concerns about job displacement due to automation, while others struggled to adapt to rapid feedback loops and shared accountability for quality. Toolchain integration issues further complicated adoption, particularly in aligning AI-powered QA tools (e.g., CodeBERT, CodeT5, XGBoost) with existing CI/CD orchestrators and monitoring systems. Interoperability gaps, inconsistent API support, and differing data formats occasionally disrupted workflows and necessitated additional middleware or custom scripts. Finally, balancing speed with thorough testing remained an ongoing tension. While automated and AI-optimized pipelines accelerated release cycles, there was a risk of omitting deep exploratory or edge-case testing in the pursuit of faster deployment. This trade-off required careful governance, dynamic test prioritization, and the preservation of manual QA checkpoints for high-risk features. Collectively, these limitations highlight that successful AI-driven QA integration is as much an organizational transformation effort as it is a technical implementation.

Conclusion and Future Work

This study demonstrated that integrating continuous QA within the DevOps pipeline, augmented by AI and machine

learning models, significantly improves software delivery performance. Deployment Success Rate increased by 16.6%, Automated Test Coverage rose by 39.1%, and Defect Leakage dropped by 64.2%, while AI-enabled optimizations reduced Mean Time to Repair and regression test runtime. These gains were achieved by embedding QA touchpoints throughout the pipeline, leveraging predictive models (e.g., XGBoost for deployment failure detection), generative AI (CodeT5 for intelligent test creation), and anomaly detection (Isolation Forest) to create a proactive, data-driven quality strategy. Beyond quantitative gains, the integration fostered a cultural shift towards shared quality ownership, continuous feedback, and faster decision-making. Future work will focus on expanding AI-driven test automation capabilities to further reduce manual intervention and improve coverage of edge cases, as well as developing self-healing pipelines capable of autonomously identifying, isolating, and correcting failures in real-time. Additional research will also explore adaptive testing strategies using reinforcement learning, enhanced interoperability between AI models and DevOps toolchains, and scalability frameworks for large, distributed engineering teams. These advancements aim to create a resilient, intelligent CI/CD ecosystem capable of sustaining rapid delivery without sacrificing quality [10-15].

References

1. Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, and M. Zhou, "CodeBERT: A Pre-Trained Model for Programming and Natural Languages," Findings of EMNLP, pp. 1536–1547, 2020.
2. Y. Wang, W. Wang, S. Joty, and S. C. H. Hoi, "CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation," in Proc. EMNLP, 2021.
3. F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in Proc. IEEE International Conference on Data Mining (ICDM), 2008, pp. 413–422.
4. N. Forsgren, J. Humble, and G. Kim (DevOps Research and Assessment — DORA), Accelerate: State of DevOps Report, DORA / Google Cloud, 2018 (annual reports available online).
5. S. Saidani, M. A. Ouni, and M. A. Mkaouer, "Predicting Continuous Integration Build Failures Using Machine-Learning and Search-Based Techniques," Information and Software Technology, 2020.
6. R. Feldt et al., "Towards Explainable Test Case Prioritisation with Learning-to-Rank Models," arXiv:2405.13786, 2024. (preprint — explores LTR approaches such as LightGBM Ranker for TCP).
7. A. A. A. Ramin et al., "A Light Bug Triage Framework for Applying Large Pre-Trained Models," Proc. ACM, 2022 — demonstrates transformer-based approaches to bug triage (BERT variants). ACM Digital Library/Wiley Online Library
8. ThoughtWorks, "Automation of Technical Tests" (Technology Radar / guidance on test automation and continuous testing best practices), ThoughtWorks, 2023–2024.
9. Prathyusha Nama et al., "Leveraging Generative AI for Automated Test Case Generation: A Framework for Enhanced Coverage and Defect Detection," Well Testing Journal, 2024.
10. D. Thakur, A. Mehra, R. Choudhary, and M. Sarker, "Generative AI in Software Engineering: Revolutionizing Test Case Generation and Validation Techniques," IRE Journals, Nov. 2023.
11. Infosys, "Generative AI for Test Automation: Revolutionizing Software Quality Assurance," White Paper, 2024.
12. ThoughtWorks, "Automation of Technical Tests," ThoughtWorks Technology Radar, 2023–2024.
13. R. Saidani, M. A. Ouni, and M. A. Mkaouer, "Predicting Continuous Integration Build Failures Using Machine-Learning and Search-Based Techniques," Information and Software Technology, 2020.
14. "Generative AI in Software Testing," Software Testing Magazine, 2024.
15. TechRadar, "Breaking silos: unifying DevOps and MLOps into a unified software supply chain," TechRadar Pro, 2025.